

PARAMETER ESTIMATION FROM TIME-SERIES DATA WITH CORRELATED ERRORS: A WAVELET-BASED METHOD AND ITS APPLICATION TO TRANSIT LIGHT CURVES

JOSHUA A. CARTER AND JOSHUA N. WINN

Department of Physics, and Kavli Institute for Astrophysics and Space Research,
 Massachusetts Institute of Technology, Cambridge, MA 02139

Accepted for publication in The Astrophysical Journal

ABSTRACT

We consider the problem of fitting a parametric model to time-series data that are afflicted by correlated noise. The noise is represented by a sum of two stationary Gaussian processes: one that is uncorrelated in time, and another that has a power spectral density varying as $1/f^\gamma$. We present an accurate and fast [$O(N)$] algorithm for parameter estimation based on computing the likelihood in a wavelet basis. The method is illustrated and tested using simulated time-series photometry of exoplanetary transits, with particular attention to estimating the midtransit time. We compare our method to two other methods that have been used in the literature, the time-averaging method and the residual-permutation method. For noise processes that obey our assumptions, the algorithm presented here gives more accurate results for midtransit times and truer estimates of their uncertainties.

Subject headings: methods: statistical — techniques: photometric — stars: planetary systems

1. INTRODUCTION

Frequently one wishes to fit a parametric model to time-series data and determine accurate values of the parameters and reliable estimates for the uncertainties in those parameters. It is important to gain a thorough understanding of the noise and develop appropriate methods for parameter estimation, especially at the research frontier, where the most interesting effects are often on the edge of detectability. Underestimating the errors leads to unjustified confidence in new results, or confusion over apparent contradictions between different data sets. Overestimating the errors inhibits potentially important discoveries.

When the errors in the data are well understood and uncorrelated, the problem of parameter estimation is relatively straightforward (see, e.g., Bevington & Robinson 2003, Gould 2003, Press et al. 2007). However, when the noise is not well-understood—and particularly when the noise exhibits correlations in time—the problem is more challenging (see, e.g., Koen & Lombard 1993, Beran 1994). Traditional methods that ignore correlations often give parameter estimates that are inaccurate and parameter errors that are underestimated. Straightforward generalization of the traditional methods is computationally intensive, with time-complexity $O(N^2)$ in the worst cases (where N is the number of data points). This makes certain analyses impractical.

Our specific concern in this paper is the analysis of time-series photometry of exoplanetary transits. During a transit, a planet passes in front of the disk of its parent star, which is evident from the slight diminution in the light received from the star. A model of a transit light curve may have many parameters, but we focus mainly on a single parameter, the midtransit time t_c , for three reasons. The first reason is the simplicity of a single-parameter model. The second reason is that t_c is a unique piece of information regarding each transit event, and as such, the accuracy cannot be improved by combining results from multiple transit observations. Instead one must make the most of single-event observations even if they are afflicted by correlated noise. The third reason is that tran-

sit timing offers a means of discovering additional planets or satellites by seeking anomalies in a sequence of transit times due to gravitational perturbations [Holman & Murray (2005), Agol et al. (2005)].¹

Beginning with the work of Pont, Zucker, & Queloz (2006), it has been widely recognized that time-correlated noise (“red noise”) is a limiting factor in the analysis of transit light curves. Many practitioners have attempted to account for correlated errors in their parameter estimation algorithms (see, e.g., Bakos et al. 2006, Gillon et al. 2006; Winn et al. 2007, 2009; Southworth 2008). Among these schemes are the “time-averaging” method, in which the effects of correlations are assessed by computing the scatter in a time-binned version of the data (Pont et al. 2006) and the “residual-permutation” method, a variant of bootstrap analysis that preserves the time ordering of the residuals (Jenkins et al. 2002).

In this paper we present an alternative method for parameter estimation in the presence of time-correlated noise, and compare it to those two previously advocated methods. The method advocated here is applicable to situations in which the noise is well described as the superposition of two stationary (time-invariant) Gaussian noise processes: one which is uncorrelated, and the other of which has a power spectral density varying as $1/f^\gamma$.

A more traditional approach to time-correlated noise is the framework of autoregressive moving average (ARMA) processes (see, e.g., Box & Jenkins 1976). The ARMA noise models can be understood as complementary to our $1/f^\gamma$ model, in that ARMA models are specified in the time domain as opposed to the frequency domain, and they are most naturally suited for modeling short-range correlations (“short-memory” processes) as opposed to long-range correlations (“long-memory” processes). Parameter estimation with ARMA models in an astronomical context has been discussed by Koen & Lombard (1993), Konig & Timmer (1997), and Timmer et al. (2000). As we will explain, our method accelerates the parameter estimation problem by taking advantage of the discrete wavelet transform. It is based on the

¹ The transit duration is also expected to vary in the presence of additional gravitating bodies; see, e.g., Kipping (2009).

fact that a the covariance matrix of a $1/f^\gamma$ noise process is nearly diagonal in a wavelet basis. As long as the actual noise is reasonably well described by such a power law, our method is attractive for its simplicity, computational speed, and ease of implementation, in addition to its grounding in the recent literature on signal processing.

The use of the wavelets in signal processing is widespread, especially for the restoration, compression, and denoising of images (see, e.g., Mallat 1999). Parameter estimation using wavelets has been considered but usually for the purpose of estimating *noise* parameters (Wornell 1996). An application of wavelets to the problem of linear regression with correlated noise was given by Fadili & Bullmore (2002). What is new in this work is the extension to an arbitrary nonlinear model, and the application to transit light curves.

This paper is organized as follows. In § 2, we review the problem of estimating model parameters from data corrupted by noise, and we review some relevant noise models. In § 3 we present the wavelet method and those aspects of wavelet theory that are needed to understand the method. In § 4, we test the method using simulated transit light curves, and compare the results to those obtained using the methods mentioned previously. In § 5 we summarize the method and the results of our tests, and suggest some possible applications and extensions of this work.

2. PARAMETER ESTIMATION WITH “COLORFUL” NOISE

Consider an experiment in which samples of an observable y_i are recorded at a sequence of times $\{t_i : i = 1, \dots, N\}$. In the context of a transit light curve, y_i is the relative brightness of the host star. We assume that the times t_i are known with negligible error. We further assume that in the absence of noise, the samples y_i would be given by a deterministic function,

$$y(t_i) = f(t_i; p_1, \dots, p_K) = f(t_i; \vec{p}), \quad (\text{no noise}) \quad (1)$$

where $\vec{p} = \{p_1, \dots, p_K\}$ is a set of K parameters that specify the function f . For an idealized transit light curve, those parameters may be the fractional loss of light δ , the total duration T , and ingress or egress duration τ , and the midtransit time t_c , in the notation of Carter et al. (2008). More realistic functions have been given by Mandel & Agol (2002) and Giménez (2007).

We further suppose that a stochastic noise process $\epsilon(t)$ has been added to the data, giving

$$y(t_i) = f(t_i; \vec{p}) + \epsilon(t_i). \quad (\text{with noise}) \quad (2)$$

As a stochastic function, $\vec{\epsilon} = \{\epsilon(t_1), \dots, \epsilon(t_N)\}$ is characterized by its joint distribution function $\mathcal{D}(\vec{\epsilon}; \vec{q})$, which in turn depends on some parameters \vec{q} and possibly also the times of observation. The goal of parameter estimation is to use the data $y(t_i)$ to calculate credible intervals for the parameters \vec{p} , often reported as best estimates \hat{p}_k and error bars $\hat{\sigma}_{p_k}$ with some quantified degree of confidence. The estimate of \vec{p} and the associated errors depend crucially on how one models the noise and how well one can estimate the relevant noise parameters \vec{q} .

In some cases one expects and observes the noise to be uncorrelated. For example, the dominant noise source may be shot noise, in which case the noise process is an uncorrelated Poisson process that in the limit of large numbers of counts is well-approximated by an uncorrelated Gaussian process,

$$\mathcal{D}(\vec{\epsilon}; \vec{q}) = \mathcal{N}(\epsilon; \sigma^2) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\epsilon_i^2}{2\sigma^2}\right), \quad (3)$$

in which case there is only one error parameter, σ , specifying the width of the distribution.

If the noise is correlated then it is characterized by a joint probability distribution that is generally a function of all the times of observation. We assume that the function is a multivariate Gaussian function, in which case the noise process is entirely characterized by the covariance matrix

$$\Sigma(t_i, t_j) = \langle \epsilon(t_i) \epsilon(t_j) \rangle. \quad (4)$$

Here, the quantity $\langle \epsilon \rangle$ is the mean of the stochastic function ϵ over an infinite number of independent realizations. We further assume that the covariance depends only on the difference in time between two samples, and not on the absolute time of either sample. In this case, the noise source is said to be stationary and is described entirely by its autocovariance $R(\tau)$ (Bracewell 1965):

$$R(\tau) \equiv \langle \epsilon(t) \epsilon(t + \tau) \rangle. \quad (5)$$

The parameter estimation problem is often cast in terms of finding the set of parameters \hat{p}_k that maximize a likelihood function. For the case of Gaussian *uncorrelated* noise the likelihood function is

$$\mathcal{L} = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \exp\left(-\frac{r_i^2}{2\hat{\sigma}^2}\right), \quad (6)$$

where r_i is the *residual* defined as $y_i - f(t_i; \vec{p})$, and $\hat{\sigma}$ is an estimate of the single noise parameter σ . Maximizing the likelihood \mathcal{L} is equivalent to minimizing the χ^2 statistic

$$\chi^2 = \sum_i^N \left(\frac{r_i}{\hat{\sigma}}\right)^2. \quad (7)$$

In transit photometry, the estimator $\hat{\sigma}$ of the noise parameter σ is usually *not* taken to be the calculated noise based on expected sources such as shot noise. This is because the actual amplitude of the noise is often greater than the calculated value due to noise sources that are unknown or at least ill-quantified. Instead, $\hat{\sigma}$ is often taken to be the standard deviation of the data obtained when the transit was not occurring, or the value for which $\chi^2 = N_{\text{dof}}$ for the best-fitting (minimum- χ^2) model. These estimates work well when the noise process is Gaussian, stationary, and uncorrelated. For the case of correlated noise, Eqn. (7) is replaced by (Gould 2003)

$$\chi^2 = \sum_{i=1}^N \sum_{j=1}^N r_i (\hat{\Sigma}^{-1})_{ij} r_j. \quad (8)$$

The case of uncorrelated noise corresponds to $\hat{\Sigma}_{ij} = \hat{\sigma}^2 \delta_{ij}$.

It is at this point where various methods for modeling correlated noise begin to diverge. One approach is to estimate $\hat{\Sigma}$ from the sample autocovariance $\hat{R}(\tau)$ of the time series, just as $\hat{\sigma}$ can be estimated from the standard deviation of the residuals in the case of uncorrelated noise. However, the calculation of χ^2 has a worst-case time-complexity of $O(N^2)$ and iterative parameter estimation techniques can be prohibitively slow. One might ameliorate the problem by truncating the covariance matrix at some maximum lag, i.e., by considering the truncated χ^2 statistic

$$\chi^2(L) = \sum_{i=1}^N \sum_{\substack{l=-L \\ 1 < i+l < N}}^L r_i (\hat{\Sigma}^{-1})_{i(i+l)} r_{i+l}, \quad (9)$$

but in the presence of long-range correlations one needs to retain many lags to obtain accurate parameter estimates. (In § 4.3, we will give an example where 50–75 lags were needed.) Alternatively, one may model the autocorrelation function and therefore the covariance matrix using an autoregressive moving-average (ARMA) model with enough terms to give a good fit to the data (see, e.g., Koen & Lombard 1993). Again, though, in the presence of long-range correlations the model covariance matrix will be non-sparse and computationally burdensome.

Pont et al. (2006) presented a useful simplification in the context of a transit search, when data are obtained on many different nights. In such cases it is reasonable to approximate the covariance matrix as block-diagonal, with different blocks corresponding to different nights. Pont et al. (2006) also gave a useful approximation for the covariance structure within each block, based on the variance in boxcar-averaged versions of the signal. Ultimately their procedure results in an equation resembling Eqn. (7) for each block, but where $\hat{\sigma}$ is the quadrature sum of σ_w (the “white noise”) and σ_r (the “red noise,” estimated from the boxcar-averaged variance). In this paper, all our examples involve a single time series with stationary noise properties, and the net effect of the Pont et al. (2006) method is to enlarge the parameter errors by a factor

$$\beta = \sqrt{1 + \left(\frac{\sigma_r}{\sigma_w}\right)^2} \quad (10)$$

relative to the case of purely white noise ($\sigma_r = 0$). We will refer to this method as the “time-averaging” method.

Another approach is to use Eqn. (7) without any modifications, but to perform the parameter optimization on a large collection of simulated data sets that are intended to have the same covariance structure as the actual data set. This is the basis of the “residual permutation” method that is also discussed further in § 4.4. As mentioned above, this method is a variant of a bootstrap analysis that takes into account time-correlated noise. More details on both the time-averaging and residual-permutation methods are given in § 4.4.

Our approach in this paper was motivated by the desire to allow for the possibility of long-range correlations, and yet to avoid the slowness of any method based on Eqn. (9) or other time-domain methods. Rather than characterizing the noise in the time domain, we characterize it by its Power Spectral Density (PSD) $\mathcal{S}(f)$ at frequency f , defined as the square of the Fourier transform of $\epsilon(t)$, or equivalently, the Fourier transform of the autocovariance $R(\tau)$. We restrict our discussion to noise sources with a PSD

$$\mathcal{S}(f) = \frac{A}{f^\gamma} \quad (11)$$

for some $A > 0$ and spectral index γ . For the special case of uncorrelated noise, $\gamma = 0$ and $\mathcal{S}(f)$ is independent of f . This type of noise has equal power density at all frequencies, which is why it is called “white noise,” in an analogy with visible light. As γ is increased, there is an increasing preponderance of low-frequency power over high-frequency power, leading to longer-range correlations in time.

Noise with a power spectrum $1/f^\gamma$ is ubiquitous in nature and in experimental science, including astrophysics (see, e.g., Press 1978). Some examples of $1/f^\gamma$ noise are shown in Fig. 1 for a selection of spectral indices. In an extension of the color analogy, $\gamma = 1$ noise is sometimes referred to as “pink noise” and $\gamma = 2$ noise as “red noise.” The latter is

also known as a Brownian process, although not because of the color brown but instead because of the Scottish botanist Robert Brown. However, as we have already noted, the term “red noise” is often used to refer to any type of low-frequency correlated noise.

Here we do not attempt to explain how $1/f^\gamma$ noise arises in a given situation. Instead we assume that the experimenter has done his or her best to understand and to reduce all sources of noise as far as possible, but despite these efforts there remains a component of $1/f^\gamma$ noise. In transit photometry these correlations often take the form of “bumps,” “wiggles,” and “ramps” in a light curve and are often attributed to differential atmospheric extinction, instrumental artifacts such as imperfect flat-fielding, and stellar granulation or other astrophysical effects. The method presented in this paper is essentially a model of the likelihood function that retains the essential information in the covariance matrix without being prohibitively expensive to compute and store. It is based on a wavelet-based description, the subject of the next section.

3. WAVELETS AND $1/f^\gamma$ NOISE

One may regard a time series with N points as a vector in an N -dimensional space that is spanned by N orthonormal unit vectors, one for each time index (the “time basis”). The computational difficulty with correlated noise is that the sample covariance matrix $\hat{\Sigma}$ is not diagonal in the time basis, nor is it necessarily close to being diagonal in realistic cases. This motivates a search for some alternative basis spanning the data space for which the covariance matrix is diagonal or nearly diagonal. For example, if the noise took the form of additive quasiperiodic signals, it would be logical to work in a Fourier basis instead of the time basis.

The mathematical result that underpins our analysis algorithm is that in the presence of $1/f^\gamma$ noise, the covariance matrix is nearly diagonal in a suitable *wavelet basis*. Before giving the details of the algorithm we will briefly review the wavelet transform. Our discussion is drawn primarily from Wornell (1996), Teolis (1998), Daubechies (1988), and Mallat (1999). Practical details and an sample implementation of the wavelet transform are given by Press et al. (2007).

A wavelet is a function that is analogous to the sine and cosine functions of the Fourier transform. Some properties that wavelets share with sines and cosines are that they are localized in frequency space, and they come in families that are related by translations and dilations. Wavelets are *unlike* sine and cosine functions in that wavelets are strongly localized in time. A wavelet basis is derived from a single “mother wavelet” $\psi(t)$, which may have a variety of functional forms and analytic properties. The individual basis functions are formed through translations and dilations of $\psi(t)$. The choice of mother wavelet depends on the specific application. We restrict our focus to dyadic orthogonal wavelet bases with basis functions

$$\psi_n^m(t) = \psi(2^m t - n) \quad (12)$$

for all integers m and n , and we further require $\psi(t)$ to have one or more vanishing moments.² In this case, the pair of equations analogous to the Fourier series and its inversion is

$$\epsilon(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \epsilon_n^m \psi_n^m(t) \quad (13)$$

² In particular it is required that the mother wavelet $\psi(t)$ has zero mean. This is a necessary and sufficient condition to ensure the invertibility of the wavelet transform.

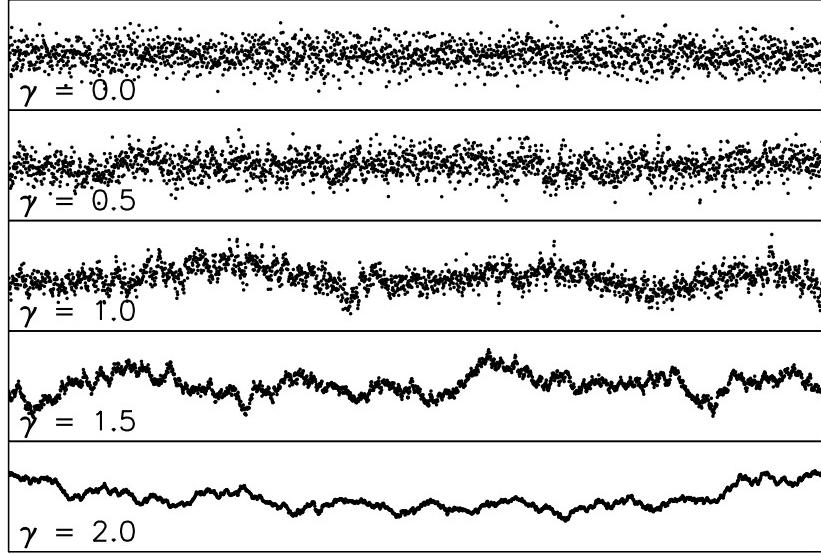


FIG. 1.— Examples of $1/f^\gamma$ noise. Uncorrelated (white) noise corresponds to $\gamma = 0$. “Pink” noise corresponds to $\gamma = 1$. “Red” noise or Brownian motion corresponds to $\gamma = 2$. These time series were generated using the wavelet-based method described in § 4.

$$\epsilon_n^m = \int_{-\infty}^{\infty} \epsilon(t) \psi_n^m(t) dt \quad (14)$$

where ϵ_n^m is referred to as the wavelet coefficient of $\epsilon(t)$ at resolution m and translation n .

3.1. The wavelet transform as a multiresolution analysis

We will see shortly that some extra terms are required in Eqn. (14) for real signals with some minimum and maximum resolution. To explain those terms it is useful to describe the wavelet transform as a multiresolution analysis, in which we consider successively higher-resolution approximations of a signal. An approximation with a resolution of 2^m samples per unit time is a member of a *resolution space* V_m . Following Wornell (1996) we impose the following conditions:

1. if $f(t) \in V_m$ then for some integer n , $f(t - 2^{-m}n) \in V_m$
2. if $f(t) \in V_m$ then $f(2t) \in V_{m+1}$.

The first condition requires that V_m contain all translations (at the resolution scale) of any of its members, and the second condition ensures that the sequence of resolutions is nested: V_m is a subset of the next finer resolution V_{m+1} . In this way, if $\epsilon_m(t) \in V_m$ is an approximation to the signal $\epsilon(t)$, then the next finer approximation $\epsilon_{m+1}(t) \in V_{m+1}$ contains all the information encoded in $\epsilon_m(t)$ plus some additional *detail* $d_m(t)$ defined as

$$d_m(t) \equiv \epsilon_{m+1}(t) - \epsilon_m(t). \quad (15)$$

We may therefore build an approximation at resolution M by starting from some coarser resolution k and adding successive detail functions:

$$\epsilon_M(t) = \epsilon_k(t) + \sum_{m=k}^M d_m(t) \quad (16)$$

The detail functions $d_m(t)$ belong to a function space $W_m(t)$, the orthogonal complement of the resolution V_m .

With these conditions and definitions, the orthogonal basis functions of W_m are the wavelet functions $\psi_n^m(t)$, obtained by translating and dilating some mother wavelet $\psi(t)$.

The orthogonal basis functions of V_m are denoted $\phi_n^m(t)$, obtained by translating and dilating a so-called “father” wavelet $\phi(t)$. Thus, the mother wavelet spawns the basis of the detail spaces, and the father wavelet spawns the basis of the resolution spaces. They have complementary characteristics, with the mother acting as a high-pass filter and the father acting as a low-pass filter.³

In Eqn. (16), the approximation $\epsilon_k(t)$ is a member of V_k , which is spanned by the functions $\phi_n^k(t)$, and $d_m(t)$ is a member of W_m , which is spanned by the functions $\psi_n^m(t)$. Thus we may rewrite Eqn. (16) as

$$\epsilon_M(t) = \sum_{n=-\infty}^{\infty} \bar{\epsilon}_n^k \phi_n^k(t) + \sum_{m=k}^M \sum_{n=-\infty}^{\infty} \epsilon_n^m \psi_n^m(t). \quad (17)$$

The *wavelet coefficients* ϵ_n^m and the *scaling coefficients* $\bar{\epsilon}_n^k$ are given by

$$\epsilon_n^m = \int_{-\infty}^{\infty} \epsilon(t) \psi_n^m(t) dt \quad (18)$$

$$\bar{\epsilon}_n^k = \int_{-\infty}^{\infty} \epsilon(t) \phi_n^k(t) dt \quad (19)$$

Eqn. (17) reduces to the wavelet-only equation (13) for the case of a continuously sampled signal $\epsilon(t)$, when we have access to all resolutions m from $-\infty$ to ∞ .⁴

There are many suitable choices for ϕ and ψ , differing in the tradeoff that must be made between smoothness and localization. The simplest choice is due to Haar (1910):

$$\phi(t) = \begin{cases} 1 & \text{if } 0 < t \leq 1 \\ 0 & \text{otherwise} \end{cases}. \quad (20)$$

$$\psi(t) = \begin{cases} 1 & \text{if } -\frac{1}{2} < t \leq 0 \\ -1 & \text{if } 0 < t \leq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases} \quad (21)$$

³ More precisely, the wavelet and scaling functions considered here are “quadrature mirror filters” (Mallat 1999).

⁴ The signal must also be bounded in order for the approximation to the signal at infinitely coarse resolution to vanish, i.e., $\lim_{k \rightarrow -\infty} \epsilon_k(t) = 0$.

The left panel of Fig. 2 shows several elements of the approximation and detail bases for a Haar multiresolution analysis. The left panels of Fig. 3 illustrate a Haar multiresolution analysis for an arbitrarily chosen signal $\epsilon(t)$, by plotting both the approximations $\epsilon_m(t)$ and details $d_m(t)$ at several resolutions m . The Haar analysis is shown for pedagogic purposes only. In practice we found it advantageous to use the more complicated fourth-order Daubechies wavelet basis, described in the next section, for which the elements and the multiresolution analysis are illustrated in the right panels of Fig. 2-3.

3.2. The Discrete Wavelet Transform

Real signals are limited in resolution, leading to finite M and k in Eqn. (17). They are also limited in time, allowing only a finite number of translations N_m at a given resolution m . Starting from Eqn. (17), we truncate the sum over n and reindex the resolution sum such that the coarsest resolution is $k = 1$, giving

$$\epsilon_M(t) = \sum_{n=1}^{N_1} \bar{\epsilon}_n^1 \phi_n^1(t) + \sum_{m=2}^M \sum_{n=1}^{N_m} \epsilon_n^m \psi_n^m(t) \quad (22)$$

where we have taken $t = 0$ to be the start of the signal. Since there is no information on timescales smaller than 2^{-M} , we need only consider $\epsilon_M(t_i)$ at a finite set of times t_i :

$$\epsilon(t_i) = \sum_{n=1}^{N_1} \bar{\epsilon}_n^1 \phi_n^1(t_i) + \sum_{m=2}^M \sum_{n=1}^{N_m} \epsilon_n^m \psi_n^m(t_i). \quad (23)$$

Eqn. (23) is the inverse of the Discrete Wavelet Transform (DWT). Unlike the continuous transform of Eqn. (13), the DWT must include the coarsest level approximation (the first term in the preceding equation) in order to preserve all the information in $\epsilon(t_i)$. For the Haar wavelet, the coarsest approximation is the mean value. For data sets with $N = n_0 2^M$ uniformly spaced samples in time, we will have access to a maximal scale M , as in Eqn. (23), with $N_m = n_0 2^{m-1}$.

A crucial point is the availability of the Fast Wavelet Transform (FWT) to perform the DWT (Mallat 1989). The FWT is a pyramidal algorithm operating on data sets of size $N = n_0 2^M$ returning $n_0(2^M - 1)$ wavelet coefficients and n_0 scaling coefficients for some $n_0 > 0$, $M > 0$. The FWT is a computationally efficient algorithm that is easily implemented (Press et al. 2007) and has $O(N)$ time-complexity (Teolis 1998).

Daubechies (1988) generalized the Haar wavelet into a larger family of wavelets, categorized according to the number of vanishing moments of the mother wavelet. The Haar wavelet has a single vanishing moment and is the first member of the family. In this work we used the most compact member (in time and frequency), $\psi = \mathcal{D}$ and $\phi = \mathcal{A}$, which is well suited to the analysis of $1/f^\gamma$ noise for $0 < \gamma < 4$ (Wornell 1996). We plot \mathcal{D}_n^m and \mathcal{A}_n^m in the time-domain for several n , m in Fig. 2, illustrating the rather unusual functional form of \mathcal{D} . The right panel of Fig. 3 demonstrates a multiresolution analysis using this basis. Press et al. (2007) provide code to implement the wavelet transform in this basis.

3.3. Wavelet transforms and $1/f^\gamma$ noise

As alluded in § 3, the wavelet transform acts as a nearly diagonalizing operator for the covariance matrix in the presence of $1/f^\gamma$ noise. The wavelet coefficients ϵ_n^m of such a noise process are zero-mean, nearly uncorrelated random variables. Specifically, the covariance between scales m, m' and translations n, n' is (Wornell 1996, p. 65)

$$\langle \epsilon_n^m \epsilon_{n'}^{m'} \rangle \approx (\sigma_r^2 2^{-\gamma m}) \delta_{m,m'} \delta_{n,n'}. \quad (24)$$

The wavelet basis is also convenient for the case in which the noise is modeled as the sum of an uncorrelated component and a correlated component,

$$\epsilon(t) = \epsilon_0(t) + \epsilon_\gamma(t), \quad (25)$$

where $\epsilon_0(t)$ is a Gaussian white noise process ($\gamma = 0$) with a single noise parameter σ_w , and $\epsilon_\gamma(t)$ has $\mathcal{S}(f) = A/f^\gamma$. In the context of transit photometry, white noise might arise from photon-counting statistics (and in cases where the detector is well-calibrated, σ_w is a known constant), while the $\gamma \neq 0$ term represents the “rumble” on many time scales due to instrumental, atmospheric, or astrophysical sources. For the noise process of Eqn. (25) the covariance between wavelet coefficients is

$$\langle \epsilon_n^m \epsilon_{n'}^{m'} \rangle \approx (\sigma_r^2 2^{-\gamma m} + \sigma_w^2) \delta_{m,m'} \delta_{n,n'}. \quad (26)$$

and the covariance between the scaling coefficients $\bar{\epsilon}_n^m$ is

$$\langle \bar{\epsilon}_n^m \bar{\epsilon}_n^{m'} \rangle \approx \sigma_r^2 2^{-\gamma m} g(\gamma) + \sigma_w^2 \quad (27)$$

where $g(\gamma)$ is a constant of order unity; for the purposes of this work $g(1) = (2 \ln 2)^{-1} \approx 0.72$ (Fadili & Bullmore 2002). Eqns. (26) and (27) are the key mathematical results that form the foundation of our algorithm. For proofs and further details, see Wornell (1996).

It should be noted that the correlations between the wavelet and scaling coefficients are small but not exactly zero. The decay rate of the correlations with the resolution index depends on the choice of wavelet basis and on the spectral index γ . By picking a wavelet basis with a higher number of vanishing moments, we hasten the decay of correlations. This is why we chose the Daubechies 4th-order basis instead of the Haar basis. In the numerical experiments described in § 4, we found the covariances to be negligible for the purposes of parameter estimation. In addition, the compactness of the Daubechies 4th-order basis reduces artifacts arising from the assumption of a periodic signal that is implicit in the FWT.

3.4. The whitening filter

Given an observation of noise $\epsilon(t)$ that is modeled as in Eqn. (25), we may estimate the $\gamma \neq 0$ component by rescaling the wavelet and scaling coefficients and filtering out the white component:

$$\epsilon_\gamma(t) = \sum_{n=1}^{N_1} \left(\frac{\sigma_r^2 2^{-\gamma} g(\gamma)}{\sigma_r^2 2^{-\gamma} g(\gamma) + \sigma_w^2} \right) \bar{\epsilon}_n^1 \phi_n^1(t) + \quad (28)$$

$$\sum_{m=2}^M \sum_{n=1}^{N_m} \left(\frac{\sigma_r^2 2^{-\gamma m}}{\sigma_r^2 2^{-\gamma m} + \sigma_w^2} \right) \epsilon_n^m \psi_n^m(t). \quad (29)$$

We may then proceed to subtract the estimate of the correlated component from the observed noise, $\epsilon_0(t) = \epsilon(t) - \epsilon_\gamma(t)$ (Wornell 1996, p. 76). In this way the FWT can be used to “whiten” the noise.

3.5. The wavelet-based likelihood

Armed with the preceding theory, we rewrite the likelihood function of Eqn. (6) in the wavelet domain. First we transform the residuals $r_i \equiv y_i - f(t_i; \vec{p})$, giving

$$r_n^m = y_n^m - f_n^m(\vec{p}) = \epsilon_{\gamma,n}^m + \epsilon_{0,n}^m \quad (30)$$

$$\bar{r}_n^1 = \bar{y}_n^1 - \bar{f}_n^1(\vec{p}) = \bar{\epsilon}_{\gamma,n}^1 + \bar{\epsilon}_{0,n}^1 \quad (31)$$

where y_n^m and $f_n^m(\vec{p})$ are the discrete wavelet coefficients of the data and the model. Likewise, \bar{y}_n^1 and $\bar{f}_n^1(\vec{p})$ are the n_0

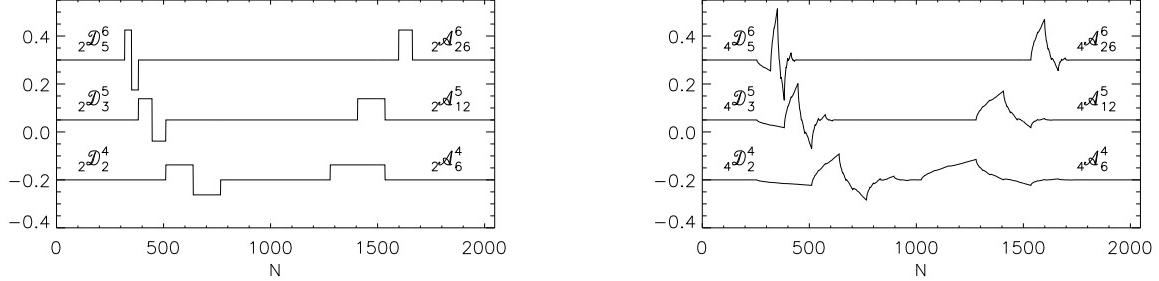


FIG. 2.— Examples of discrete wavelet and scaling functions, for $N = 2048$. *Left*.—Haar wavelets and the corresponding father wavelets, also known as 2nd-order Daubechies orthonormal wavelets or ${}_2\mathcal{D}_n^m$ and ${}_2\mathcal{A}_n^m$. *Right*.—4th-order Daubechies orthonormal wavelets, or ${}_4\mathcal{D}_n^m$ and ${}_4\mathcal{A}_n^m$.

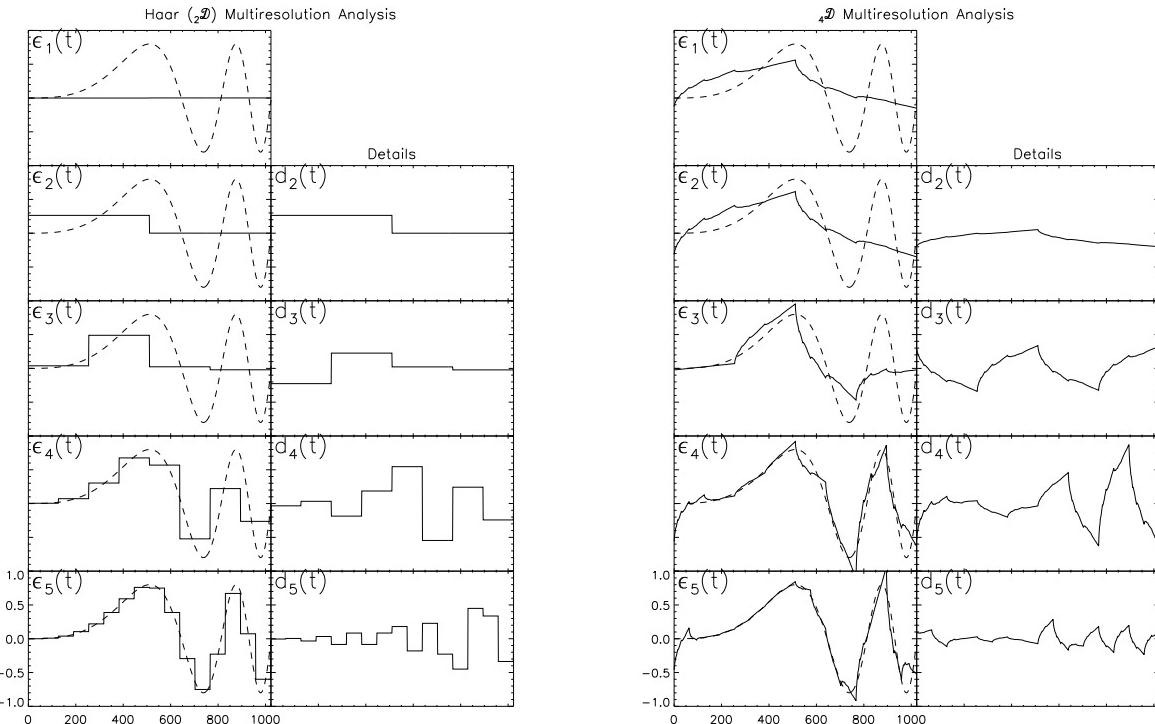


FIG. 3.— Illustration of a multiresolution analysis, for the function $e(t) = \sin[4\pi(t/1024)^3]$ (dashed line). Plotted are the approximations $\epsilon_m(t)$ to the function at successive resolutions, along with the detail functions $d_m(t)$. *Left*.—Using the Haar wavelet basis. *Right*.—Using the 4th-order Daubechies wavelet basis.

scaling coefficients of the data and the model. Given the diagonal covariance matrix shown in Eqns. (26) and (27), the likelihood \mathcal{L} is a product of Gaussian functions at each scale m and translation n :

$$\begin{aligned} \mathcal{L} &= \left\{ \prod_{m=2}^M \prod_{n=1}^{n_0 2^{m-1}} \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp \left[-\frac{(r_n^m)^2}{2\sigma_W^2} \right] \right\} \\ &\times \left\{ \prod_{n=1}^{n_0} \frac{1}{\sqrt{2\pi\sigma_S^2}} \exp \left[-\frac{(\bar{r}_n^1)^2}{2\sigma_S^2} \right] \right\} \end{aligned} \quad (32)$$

where

$$\sigma_W^2 = \sigma_r^2 2^{-\gamma m} + \sigma_w^2 \quad (33)$$

$$\sigma_S^2 = \sigma_r^2 2^{-\gamma} g(\gamma) + \sigma_w^2 \quad (34)$$

are the variances of the wavelet and scaling coefficients respectively. For a data set with N points, calculating the likelihood function of Eqn. (32) requires multiplying N Gaussian functions. The additional step of computing the FWT of the

residuals prior to computing \mathcal{L} adds $O(N)$ operations. Thus, the entire calculation has a time-complexity $O(N)$.

For this calculation we must have estimators of the three noise parameters γ , σ_r , and σ_w . These may be estimated separately from the model parameters \vec{p} , or simultaneously with the model parameters. For example, in transit photometry, the data obtained outside of the transit may be used to estimate the noise parameters, which are then used in Eqn. (32) to estimate the model parameters. Or, in a single step we could maximize Eqn. (32) with respect to all of γ , σ_r , σ_w and \vec{p} . Fitting for both noise and transit parameters simultaneously is potentially problematic, because some of the correlated noise may be “absorbed” into the choices of the transit parameters, i.e., the errors in the noise parameters and transit parameters are themselves correlated. This may cause the noise level and the parameter uncertainties to be underestimated. Unfortunately, there are many instances when one does not have enough out-of-transit data for the strict separation of transit and noise parameters to be feasible.

In practice the optimization can be accomplished with an

iterative routine [such as AMOEBA, Powell’s method, or a conjugate-gradient method; see Press et al. (2007)]. Confidence intervals can then be defined by the contours of constant likelihood. Alternatively one can use a Monte Carlo Markov Chain [MCMC; see, e.g., Gregory (2005)], in which case the jump-transition likelihood would be given by Eqn. (32). The advantages of the MCMC method have led to its adoption by many investigators (see, e.g., Holman et al. 2006, Burke et al. 2007, Collier Cameron et al. 2007). For that method, computational speed is often a limiting factor, as a typical MCMC analysis involves several million calculations of the likelihood function.

3.6. Some practical considerations

Some aspects of real data do not fit perfectly into the requirements of the DWT. The time sampling of the data should be approximately uniform, so that the resolution scales of the multiresolution analysis accurately reflect physical timescales. This is usually the case for time-series photometric data. Gaps in a time series can be fixed by applying the DWT to each uninterrupted data segment, or by filling in the missing elements of the residual series with zeros.

The FWT expects the number of data points to be an integral multiple of some integral power of two. When this is not the case, the time series may be truncated to the nearest such boundary; or it may be extended using a periodic boundary condition, mirror reflection, or zero-padding. In the numerical experiments described below, we found that zero-padding has negligible effects on the calculation of likelihood ratios and parameter estimation.

The FWT generally assumes a periodic boundary condition for simplicity of computation. A side effect of this simplification is that information at the beginning and end of a time series are artificially associated in the wavelet transform. This is one reason why we chose the 4th-order Daubechies-class wavelet basis, which is well localized in time, and does not significantly couple the beginning and the end of the time series except on the coarsest scales.

4. NUMERICAL EXPERIMENTS WITH TRANSIT LIGHT CURVES

We performed many numerical experiments to illustrate and test the wavelet method. These experiments involved estimating the parameters of simulated transit light curves. We also compared the wavelet analysis to a “white” analysis, by which we mean a method that assumes the errors to be uncorrelated, and to two other analysis methods drawn from the literature. Because we used simulated transit light curves with known noise and transit parameters, the “truth” was known precisely, allowing both the absolute and relative merits of the methods to be evaluated.

4.1. Estimating the midtransit time: Known noise parameters

In this section we consider the case in which the noise parameters γ , σ_r , and σ_w are known with negligible error. We have in mind a situation in which a long series of out-of-transit data are available, with stationary noise properties.

We generated transit light curves with known transit parameters \vec{p} , contaminated by an additive combination of a white and a correlated ($1/f^\gamma$) noise source. Then we used an MCMC method to estimate the transit parameters and their 68.3% confidence limits. (The technique for generating noise and the MCMC method are described in detail below.) For each realization of a simulated light curve, we estimated tran-

sit parameters using the likelihood defined either by Eqn. (6) for the white analysis, or Eqn. (32) for the wavelet analysis.

For a given parameter p_k , the estimator \hat{p}_k was taken to be the median of the values in the Markov chain and $\hat{\sigma}_{p_k}$ was taken to be the standard deviation of those values. To assess the results, we considered the “number-of-sigma” statistic

$$\mathcal{N} \equiv (\hat{p}_k - p_k) / \hat{\sigma}_{p_k}. \quad (35)$$

In words, \mathcal{N} is the number of standard deviations separating the parameter estimate \hat{p}_k from the true value p_k . If the error in p_k is Gaussian, then a perfect analysis method should yield results for \mathcal{N} with an expectation value of 0 and variance of 1. If we find that the variance of \mathcal{N} is greater than one, then we have underestimated the error in \hat{p}_k and we may attribute too much significance to the result. On the other hand, if the variance of \mathcal{N} is smaller than one, then we have overestimated σ_{p_k} and we may miss a significant discovery. If we find that the mean of \mathcal{N} is nonzero then the method is biased.

For now, we consider only the single parameter t_c , the time of midtransit. The t_c parameter is convenient for this analysis as it is nearly decoupled from the other transit parameters (Carter et al. 2008). Furthermore, as mentioned in the introduction, the measurement of the midtransit time cannot be improved by observing other transit events, and variations in the transit interval are possible signs of additional gravitating bodies in a planetary system.

The noise was synthesized as follows. First, we generated a sequence of $N = 1024$ independent random variables obeying the variance conditions from Eqns. (26) and (27) for 1023 wavelet coefficients over 9 scales and a single scaling coefficient at the coarsest resolution scale. We then performed the inverse FWT of this sequence to generate our noise signal. In this way, we could select exact values for γ , σ_r , and σ_w . We also needed to find the single parameter σ for the white-noise analysis; it is not simply related to the parameters γ , σ_r , and σ_w . In practice, we found σ by calculating the median sample variance among 10^4 unique realizations of a noise source with fixed parameters γ , σ_r , and σ_w .

For the transit model, we used the analytic formulas of Mandel & Agol (2002), with a planet-to-star ratio of $R_p/R_\star = 0.15$, a normalized orbital distance of $a/R_\star = 10$, and an orbital inclination of $i = 90^\circ$, as appropriate for a gas giant planet in a close-in orbit around a K star. These correspond to a fractional loss of light $\delta = 0.0225$, duration $T = 1.68$ hr, and partial duration $\tau = 0.152$ hr. We did not include the effect of limb darkening, as it would increase the computation time and has little influence on the determination of t_c (Carter et al. 2009). Each simulated light curve spanned 3 hr centered on the midtransit time, with a time sampling of 11 s, giving 1024 uniformly spaced samples. A noise-free light curve is shown in Fig. 4.

For the noise model, we chose $\sigma_w = 1.35 \times 10^{-3}$ and $\gamma = 1$, and tried different choices for σ_r . We denote by α the ratio of the rms values of the correlated noise component and the white noise component.⁵ The example in Fig. 4 has $\alpha = 1/3$. As α is increased from zero, the correlated component becomes more important, as is evident in the simulated data plotted in Fig. 5. Our choice of σ_w corresponds to a precision of 5.8×10^{-4} per minute-equivalent sample, and was inspired by the recent work by Johnson et al. (2009) and Winn

⁵ We note that although σ_w is the rms of the white noise component, σ_r is generally not the rms of the correlated component. The notation is unfortunate, but follows that of Wornell (1996).

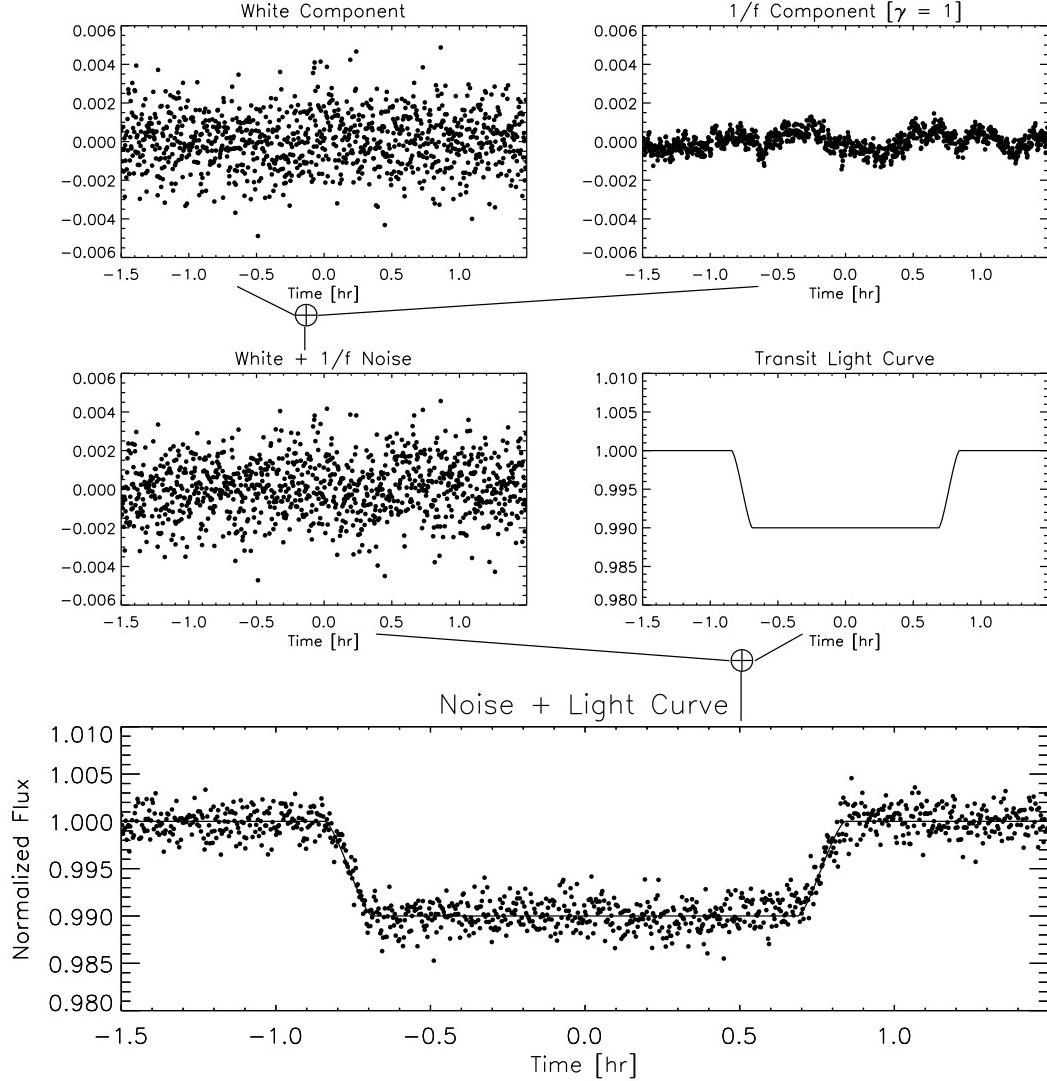


FIG. 4.—Constructing a simulated transit light curve with correlated noise. The total noise is the sum of uncorrelated Gaussian noise with standard deviation σ_w (upper left panel) and correlated noise with a power spectral density $S(f) \propto 1/f$ and an rms equal to $\sigma_w/3$ (upper right panel). The total noise (middle left panel) is added to an idealized transit model (middle right panel) to produce the simulated data (bottom panel).

et al. (2009), which achieved precisions of 5.4×10^{-4} and 4.0×10^{-4} per minute-equivalent sample, respectively. Based on our survey of the literature and our experience with the Transit Light Curve project (Holman et al. 2006, Winn et al. 2007), we submit that all of the examples shown in Fig. 5 are “realistic” in the sense that the bumps, wiggles, and ramps resemble features in actual light curves, depending on the instrument, observing site, weather conditions, and target star.

For a given choice of α , we made 10,000 realizations of the simulated transit light curve with $1/f$ noise. We then constructed two Monte Carlo Markov Chains for t_c starting at the true value of $t_c = 0$. One chain was for the white analysis, with a jump-transition likelihood given by Eqn. (6). The other chain was for the wavelet analysis, using Eqn. (32) instead. Both chains used the Metropolis-Hastings jump condition, and employed perturbation sizes such that $\approx 40\%$ of jumps were accepted. Initial numerical experiments showed that the autocorrelation of a given Markov chain for t_c is sharply peaked at zero lag, with the autocorrelation dropping below 0.2 at lag-one. This ensured good convergence with

chain lengths of 500 (Tegmark et al. 2004). Chain histograms were also inspected visually to verify that the distribution was smooth. We recorded the median \hat{t}_c and standard deviation $\hat{\sigma}_{t_c}$ for each chain and constructed the statistic \mathcal{N} for each separate analysis (white or wavelet). Finally, we found the median and standard deviation of \mathcal{N} over all 10,000 noise realizations.

Fig. 6 shows the resulting distributions of \mathcal{N} , for the particular case $\alpha = 1/3$. Table 1 gives a collection of results for the choices $\alpha = 0, 1/3, 2/3$, and 1. The mean of \mathcal{N} is zero for both the white and wavelet analyses: neither method is biased. This is expected, because all noise sources were described by zero-mean Gaussian distributions. However, the widths of the distributions of \mathcal{N} show that the white analysis underestimates the error in t_c . For a transit light curve constructed with equal parts white and $1/f$ noise ($\alpha = 1$), the white analysis gave an estimate of t_c that differs from the true value by more than 1σ nearly 80% of the time. The factor by which the white analysis underestimates the error in t_c appears to increase linearly with α . In contrast, for all values of α , the wavelet analysis maintains a unit variance in \mathcal{N} , as desired.

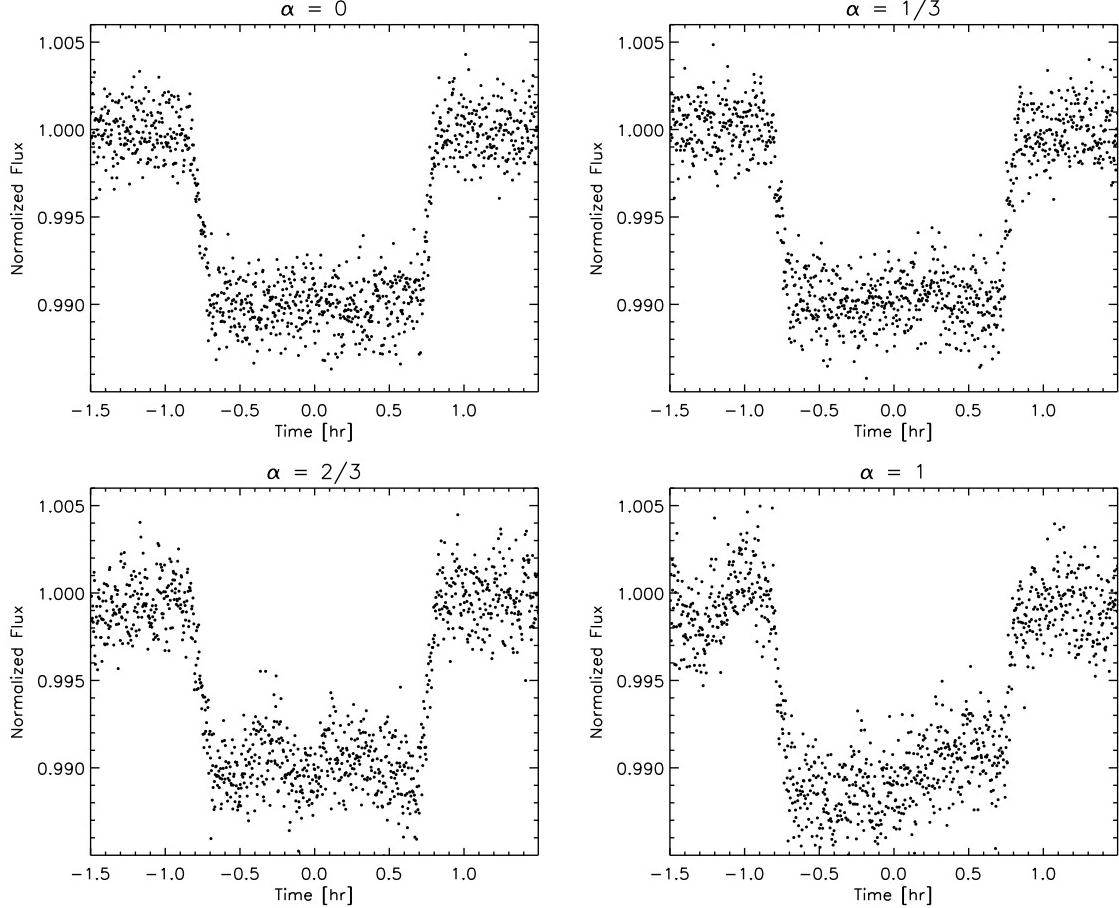


FIG. 5.— Examples of simulated transit light curves with different ratios $\alpha = \text{rms}_r/\text{rms}_w$ between the rms values of the correlated noise component and white noise component.

TABLE 1
ESTIMATES OF MID-TRANSIT TIME, t_c , FROM DATA WITH KNOWN NOISE PROPERTIES

Method	α	$\langle \hat{\sigma}_{t_c} \rangle$ [sec]	$\langle \mathcal{N} \rangle$	$\sigma_{\mathcal{N}}$	$\text{prob}(\mathcal{N} > 1)$	prob(best)^a
White	0	4.1	+0.004	0.95	29%	50%
	1/3	4.3	-0.005	1.93	61%	39%
	2/3	5.0	+0.005	3.04	75%	35%
	1	5.9	-0.036	3.82	79%	34%
Wavelet	0	4.0	+0.005	0.95	29%	50%
	1/3	7.2	-0.004	0.93	28%	61%
	2/3	11.5	-0.004	0.94	28%	65%
	1	16.0	-0.001	0.95	29%	66%

^a The probability that the analysis method (white or wavelet) returns an estimate of t_c that is closer to the true value than the other method.

The success of the wavelet method is partially attributed to the larger (and more appropriate) error intervals that it returns for \hat{t}_c . It is also partly attributable to an improvement in the accuracy of \hat{t}_c itself: the wavelet method tends to produce \hat{t}_c values that are closer to the true t_c . This is shown in the final column in Table (1), where we report the percentage of cases in which the analysis method (white or wavelet) produces an estimate of t_c that is closer to the truth. For $\alpha = 1$ the wavelet analysis gives more accurate results 66% of the time.

4.2. Estimating the midtransit time: Unknown noise parameters

In this section we consider the case in which the noise parameters are not known in advance. Instead the noise parameters must be estimated based on the data. We did this by including the noise parameters as adjustable parameters in the Markov chains. In principle this could be done for all three noise parameters γ , σ_r , and σ_w , but for most of the experiments presented here we restricted the problem to the case $\gamma = 1$. This may be a reasonable simplification, given the preponderance of natural noise sources with $\gamma = 1$ (Press 1978). Some experiments involving noise with $\gamma \neq 1$ are described at the end of this section.

We also synthesized the noise with a non-wavelet technique, to avoid “stacking the deck” in favor of the wavelet method. We generated the noise in the frequency domain, as follows. We specified the amplitudes of the Fourier coefficients using the assumed functional form of the power spectral density [$S(f) \propto 1/f$], and drew the phases from a uniform distribution between $-\pi$ and π . The correlated noise in the time domain was found by performing an inverse Fast Fourier Transform. We rescaled the noise such that the rms was α times the specified σ_w . The normally-distributed white noise was then added to the correlated noise to create the total noise. This in turn was added to the idealized transit model.

For each choice of α , we made 10,000 simulated transit light curves and analyzed them with the MCMC method described previously. For the white analysis, the mid-transit

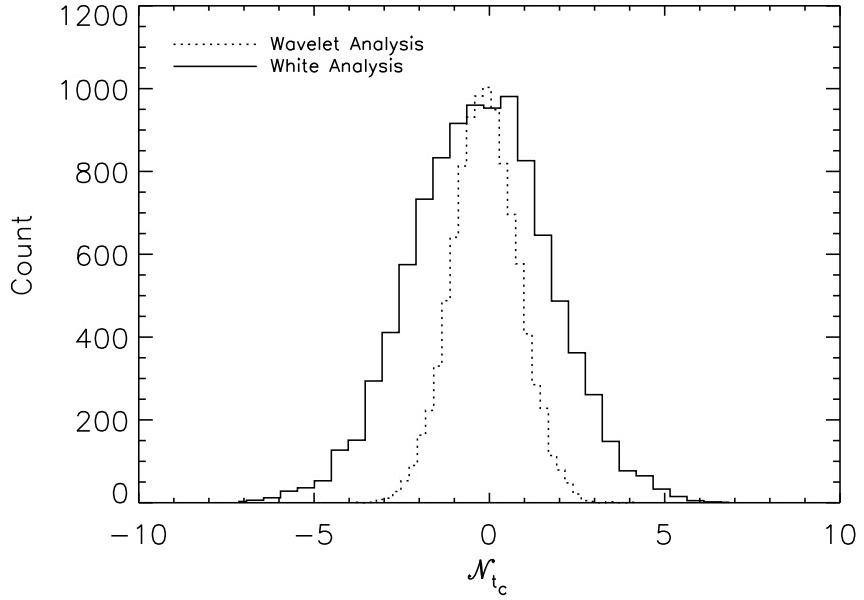


FIG. 6.— Histograms of the number-of-sigma statistic \mathcal{N} for the midtransit time t_c . Each distribution shows the probability of estimating a value for t_c that differs by $\mathcal{N}\sigma$ from the true value. The simulated data were created by adding an idealized transit model to a noise source that is the sum of uncorrelated noise and $1/f$ noise with equal variances ($\alpha = 1$; see the text).

time t_c and the single noise parameter σ were estimated using the likelihood defined via Eqn. (6). For the wavelet analysis we estimated t_c and the two noise parameters σ_r and σ_w using the likelihood defined in Eqn. (32).

Table 3 gives the resulting statistics from this experiment, in the same form as were given in Table 1 for the case of known noise parameters. (This table also includes some results from § 4.4, which examines two other methods for coping with correlated noise.) Again we find that the wavelet method produces a distribution of \mathcal{N} with unit variance, regardless of α ; and again, we find that the white analysis underestimates the error in t_c . In this case the degree of error underestimation is less severe, a consequence of the additional freedom in the noise model to estimate σ from the data. The wavelet method also gives more accurate estimates of t_c than the white method, although the contrast between the two methods is smaller than it was with for the case of known noise parameters.

Our numerical results must be understood to be illustrative, and not universal. They are specific to our choices for the noise parameters and transit parameters. Via further numerical experiments, we found that the width of \mathcal{N} in the white analysis is independent of σ_w , but it does depend on the time sampling. In particular, the width grows larger as the time sampling becomes finer (see Table 2). This can be understood as a consequence of the long-range correlations. The white analysis assumes that the increased number of data points will lead to enhanced precision, whereas in reality, the correlations negate the benefit of finer time sampling.

Table 4 gives the results of additional experiments with $\gamma \neq 1$. In those cases we created simulated noise with $\gamma \neq 1$ but in the course of the analysis we assumed $\gamma = 1$. The correlated noise fraction was set to $\alpha = 1/2$ for these tests. The results show that even when γ is falsely assumed to be unity, the wavelet analysis still produces better estimates of t_c and more reliable error bars than the white analysis.

TABLE 2
EFFECT OF TIME SAMPLING ON THE WHITE ANALYSIS

N^a	Cadence [sec]	$\sigma_{\mathcal{N}}$
256	42.2	1.72
512	21.1	2.04
1024	10.5	2.69
2048	5.27	3.49
4096	2.63	4.39

^a The number of samples in a 3 hr interval.

4.3. Runtime analysis of the time-domain method

Having established the superiority of the wavelet method over the white method, we wish to show that the wavelet method is also preferable to the more straightforward approach of computing the likelihood function in the time domain with a non-diagonal covariance matrix. The likelihood in this case is given by Eqn. (8).

The time-domain calculation and the use of the covariance matrix raised two questions. First, how well can we estimate the autocovariance $R(\tau)$ from a single time series? Second, how much content of the resulting covariance matrix needs to be retained in the likelihood calculation for reliable parameter estimation? The answer to the first question depends on whether we wish to utilize the sample autocorrelation as the estimator of $R(\tau)$ or instead use a parametric model (such as an ARMA model) for the autocorrelation. In either case, our ability to estimate the autocorrelation improves with number of data samples contributing to its calculation. The second question is important because retaining the full covariance matrix would cause the computation time to scale as $O(N^2)$ and in many cases the analysis would be prohibitively slow. The second question may be reframed as: what is the minimum number of lags L that needs to be considered in computing the truncated χ^2 of Eqn. (9), in order to give unit variance in the number-of-sigma statistic for each model parameter?

TABLE 3
ESTIMATES OF t_c FROM DATA WITH UNKNOWN NOISE PROPERTIES

Method	α	$\langle \hat{\sigma}_{t_c} \rangle$ [sec]	$\langle \mathcal{N} \rangle$	$\sigma_{\mathcal{N}}$	$prob(\mathcal{N} > 1)$	$prob(\text{better})^a$
White	0	4.0	-0.011	0.97	31%	—
	1/3	4.2	+0.010	1.70	57%	—
	2/3	4.9	+0.012	2.69	73%	—
	1	5.8	+0.023	3.28	78%	—
Wavelet	0	4.5	-0.009	0.90	26%	50%
	1/3	6.9	-0.003	1.03	33%	56%
	2/3	11.2	-0.005	1.07	35%	57%
	1	15.7	-0.007	1.09	36%	57%
Time-averaging	0	4.4	-0.006	0.88	26%	50%
	1/3	6.8	+0.009	1.15	36%	50%
	2/3	11.6	-0.012	1.24	40%	50%
	1	17.6	+0.007	1.21	38%	50%
Residual-permutation	0	3.5	-0.012	1.16	37%	50%
	1/3	6.6	+0.013	1.24	37%	50%
	2/3	11.8	-0.014	1.28	38%	49%
	1	17.3	+0.008	1.30	38%	48%

^a The probability that the analysis method returns an estimate of t_c that is closer to the true value than the white analysis.

TABLE 4
ESTIMATES OF t_c FROM DATA WITH UNKNOWN NOISE PROPERTIES

Method	γ^a	$\langle \hat{\sigma}_{t_c} \rangle$ [sec]	$\langle \mathcal{N} \rangle$	$\sigma_{\mathcal{N}}$	$prob(\mathcal{N} > 1)$	$prob(\text{best})^b$
White	0.5	4.5	-0.025	1.34	47%	50%
	1.5	4.6	+0.020	3.10	77%	32%
Wavelet	0.5	6.7	-0.021	0.97	30%	50%
	1.5	6.9	+0.002	1.17	39%	68%

^a The spectral exponent of the Power Spectral Density, $S(f) \propto 1/f^\gamma$. ^b The probability that the analysis method (white or wavelet) returns an estimate of t_c that is closer to the true value than the other method.

The time-complexity of the truncated likelihood calculation is $O(NL)$. If $L \lesssim 5$ then the time-domain method and the wavelet method may have comparable computational time-complexity, while for larger L the wavelet method would offer significant advantage.

We addressed these questions by repeating the experiments of the previous sections using a likelihood function based on the truncated χ^2 statistic. We assumed that the parameters of the noise model were known, as in § 4.1. The noise was synthesized in the wavelet domain, with $\gamma = 1$, $\sigma_w = 0.00135$, and α set equal to 1/3 or 2/3. The parameters of the transit model and the time series were the same as in § 4.1. We calculated the “exact” autocovariance function $R(l)$ at integer lag l for a given α by averaging sample autocovariances over 50,000 noise realizations. Fig. 7 plots the autocorrelation [$R(l)/R(0)$] as a function of lag for $\alpha = 1/3$, 2/3. We constructed the stationary covariance $\Sigma_{ij} = R(|i - j|)$ and computed its inverse (Σ^{-1})_{ij} for use in Eqn. (9).

Then we used the MCMC method to find estimates and errors for the time of midtransit, and calculated the number-of-sigma statistic \mathcal{N} as defined in Eqn. (35). In particular, for each simulated transit light curve, we created a Markov chain of 1,000 links for t_c , using $\chi^2(L)$ in the jump-transition like-

lihood. We estimated t_c and σ_{t_c} , and calculated \mathcal{N} . We did this for 5,000 realizations and determined $\sigma_{\mathcal{N}}$, the variance in \mathcal{N} , across this sample. We repeated this process for different choices of the maximum lag L . Fig. (8) shows the dependence of $\sigma_{\mathcal{N}}$ upon the maximum lag L .

The time-domain method works fine, in the sense that when enough non-diagonal elements in the covariance matrix are retained, the parameter estimation is successful. We find that $\sigma_{\mathcal{N}}$ approaches unity as $L^{-\beta}$ with $\beta = 0.15$, 0.25 for $\alpha = 1/3$, 2/3, respectively. However, to match the reliability of the wavelet method, a large number of lags must be retained. To reach $\sigma_{\mathcal{N}} = 1.05$, we need $L \approx 50$ for $\alpha = 1/3$ or $L \approx 75$ for $\alpha = 2/3$. In our implementation, the calculation based on the truncated covariance matrix [Eqn. (9)] took 30–40 times longer than the calculation based on the wavelet likelihood [Eqn. (32)].

This order-of-magnitude penalty in runtime is bad enough, but the real situation may be even worse, because one usually has access to a single noisy estimate of the autocovariance matrix. Or, if one is using an ARMA model, the estimated parameters of the model might be subject to considerable uncertainty as compared to the “exact” autocovariance employed in our numerical experiments. If it is desired to determine the noise parameters simultaneously with the other model parameters, then there is a further penalty associated with inverting the covariance matrix at each step of the calculation for use in Eqn. (9), although it may be possible to circumvent that particular problem by modeling the inverse-covariance matrix directly.

4.4. Comparison with other methods

In this section we compare the results of the wavelet method to two methods for coping with correlated noise that are drawn from the recent literature on transit photometry. The first of these two methods is the “time averaging” method that was propounded by Pont et al. (2006) and used in various forms by Bakos et al. (2006), Gillon et al. (2006), Winn et al. (2007, 2008, 2009), Gibson et al. (2008), and others. In

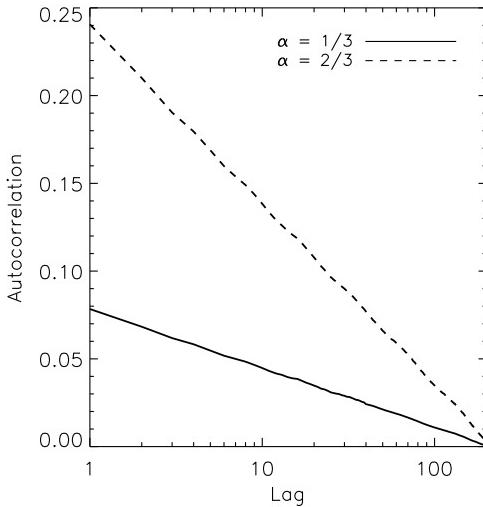


FIG. 7.— Autocorrelation functions of correlated noise. The noise was computed as the sum of white noise with $\sigma_w = 0.00135$ and $1/f$ noise with an rms equal to $\alpha\sigma_w$, for $\alpha = 1/3$ or $2/3$.

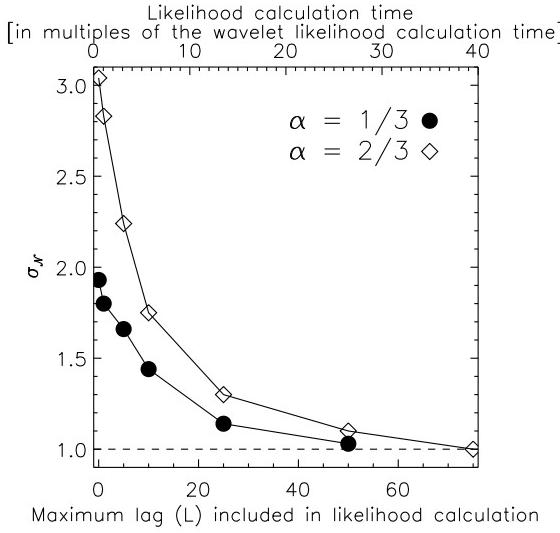


FIG. 8.— Accuracy of the truncated time-domain likelihood in estimating midtransit times. Plotted is the variance in the number-of-sigma statistic σ_N for the midtransit time t_c , as a function of the maximum lag in the truncated series. The estimates of t_c were found using the truncated likelihood given in Eqn. (9).

one implementation, the basic idea is to calculate the sample variance of unbinned residuals, $\hat{\sigma}_1^2$, and also the sample variance of the time-averaged residuals, $\hat{\sigma}_n^2$, where every n points have been averaged (creating m time bins). In the absence of correlated noise, we expect

$$\hat{\sigma}_n^2 = \frac{\hat{\sigma}_1^2}{n} \left(\frac{m}{m-1} \right). \quad (36)$$

In the presence of correlated noise, $\hat{\sigma}_n^2$ differs from this expectation by a factor $\hat{\beta}_n^2$. The estimator $\hat{\beta}$ is then found by averaging $\hat{\beta}_n$ over a range Δn corresponding to time scales that are judged to be most important. In the case of transit photometry, the duration of ingress or egress is the most relevant time scale (corresponding to averaging time scales on the order of tens of minutes, in our example light curve). A white analysis is then performed, using the noise parameter

$\sigma = \beta\sigma_1$ instead of σ_1 . This causes the parameter errors $\hat{\sigma}_{pk}$ to increase by β but does not change the parameter estimates \hat{p}_k themselves.⁶

A second method is the “residual permutation” method that has been used by Jenkins et al. (2002), Moutou et al. (2004), Southworth (2008), Bean et al. (2008), Winn et al. (2008), and others. This method is a variant of a bootstrap analysis, in which the posterior probability distribution for the parameters is based on the collection of results of minimizing χ^2 (assuming white noise) for a large number of synthetic data sets. In the traditional bootstrap analysis the synthetic data sets are produced by scrambling the residuals and adding them to a model light curve, or by drawing data points at random (with replacement) to make a simulated data set with the same number of points as the actual data set. In the residual permutation method, the synthetic data sets are built by performing a cyclic permutation of the time indices of the residuals, and then adding them to the model light curve. In this way, the synthetic data sets have the same bumps, wiggles, and ramps as the actual data, but they are translated in time. The parameter errors are given by the widths of the distributions in the parameters that are estimated from all the different realizations of the synthetic data, and they are usually larger than the parameter errors returned by a purely white analysis.

As before, we limited the scope of the comparison to the estimation of t_c and its uncertainty. We created 5,000 realizations of a noise source with $\gamma = 1$ and a given value of α (either 0, $1/3$, $2/3$, or 1). We used each of the two approximate methods (time-averaging and residual-permutation) to calculate $\hat{\beta}$ and its uncertainty based on each of the 5,000 noise realizations. Then we found the median and standard deviation of $\hat{\beta}/\beta$ over all 5,000 realizations. Table (3) presents the results of this experiment.

Both methods, time-averaging and residual-permutation, gave more reliable uncertainties than the white method. However they both underestimated the true uncertainties by approximately 15–30%. Furthermore, neither method provided more accurate estimates of t_c than did the white method. For the time-averaging method as we have implemented it, this result is not surprising, for that method differs from the white method only in the inflation of the error bars by some factor β . The parameter values that maximize the likelihood function were unchanged.

4.5. Alternative noise models

We have shown the wavelet method to work well in the presence of $1/f^\gamma$ noise. Although this family of noise processes encompasses a wide range of possibilities, the universe of possible correlated noise processes is much larger. In this section we test the wavelet method using simulated data that has correlated noise of a completely different character. In particular, we focus on a process with exclusively short-term correlations, described by one of the aforementioned autoregressive moving-average (ARMA) class of parametric noise models. In this way we test our method on a noise process that is complementary to the longer-range correlations present in $1/f^\gamma$ noise, and we also make contact between our method and the large body of statistical literature on ARMA models.

⁶ Alternatively one may assign an error to each data point equal to the quadrature sum of the measurement error and an extra term σ_r (Pont et al. 2006). For cases in which the errors in the data points are all equal or nearly equal, these methods are equivalent. When the errors are not all the same, it is more appropriate to use the quadrature-sum approach of Pont et al. (2006). In this paper all our examples involve homogeneous errors.

For $1/f^\gamma$ noise we have shown that time-domain parameter estimation techniques are slow. However, if the noise has exclusively short-range correlations, the autocorrelation function will decay with lag more rapidly than a power law, and the truncated- χ^2 likelihood [Eqn. (9)] may become computationally efficient. ARMA models provide a convenient analytic framework for parameterizing such processes. For a detailed review of ARMA models and their use in statistical inference, see Box & Jenkins (1976). Applications of ARMA models to astrophysical problems have been described by in Koen & Lombard (1993), Konig & Timmer (1997) and Timmer et al. (2000).

To see how the wavelet method performs on data with short-range correlations we constructed synthetic transit data in which the noise is described by a single-parameter autoregressive [AR(1; ψ)] model. An AR(1; ψ) process $\epsilon(t_i)$ is defined by the recursive relation

$$\epsilon(t_i) = \eta(t_i) + \psi\epsilon(t_{i-1}) \quad (37)$$

where $\eta(t_i)$ is an uncorrelated Gaussian process with width parameter σ and ψ is the sole autoregressive parameter. The autocorrelation $\gamma(l)$ for an AR(1; ψ) process is

$$\gamma(l) = \frac{\sigma^2}{1-\psi^2}\psi^l. \quad (38)$$

An AR(1; ψ) process is stationary so long as $0 < \psi < 1$ (Box & Jenkins 1976). The decay length of the autocorrelation function grows as ψ is increased from zero to one. Figure (9) plots the autocorrelation function of a process that is an additive combination of an AR(1; $\psi = 0.95$) process and a white noise process. The noise in our synthetic transit light curves was the sum of this AR(1; $\psi = 0.95$) process, and white noise, with $\alpha = 1/2$ (see Fig 9). With these choices, the white method underestimates the error in t_c , while at the same time the synthetic data look realistic.

We proceeded with the MCMC method as described previously to estimate the time of mid-transit. All four methods assessed in the previous section were included in this analysis, for comparison. Table 5 gives the results. The wavelet method produces more reliable error estimates than the white method. However, the wavelet method no longer stands out as superior to the time-averaging method or the residual-permutation method; all three of these methods give similar results. This illustrates the broader point that using any of these methods is much better than ignoring the noise correlations. The results also show that although the wavelet method is specifically tuned to deal with $1/f^\gamma$ noise, it is still useful in the presence of noise with shorter-range correlations.

It is beyond the scope of this paper to test the applicability of the wavelet method on more general ARMA processes. Instead we suggest the following approach when confronted with real data [see also Beran (1994)]. Calculate the sample autocorrelation, and power spectral density, based on the out-of-transit data or the residuals to an optimized transit model. For stationary processes these two indicators are related as described in § 2. Short-memory, ARMA-like processes can be identified by large autocorrelations at small lags or by finite power spectral density at zero frequency. Long-memory processes ($1/f^\gamma$) can be identified by possibly small but non-vanishing autocorrelation at longer lags. Processes with short-range correlations could be analyzed with an ARMA model of the covariance matrix [see Box & Jenkins (1976)], or the truncated-lag covariance matrix, although a wavelet-based analysis may be sufficient as well. Long-memory processes

are best analyzed with the wavelet method as described in this paper.

It should also be noted that extensions of ARMA models have been developed to mimic long-memory, $1/f^\gamma$ processes. In particular, fractional autoregressive integrated moving-average models (ARFIMA) describe “nearly” $1/f^\gamma$ stationary processes, according to the criterion described by Beran (1994). As is the case with ARMA models, ARFIMA models enjoy analytic forms for the likelihood in the time-domain. Alas, as noted by Wornell (1996) and Beran (1994), the straightforward calculation of this likelihood is computationally expensive and potentially unstable. For $1/f^\gamma$ processes, the wavelet method is probably a better choice than any time-domain method for calculating the likelihood.

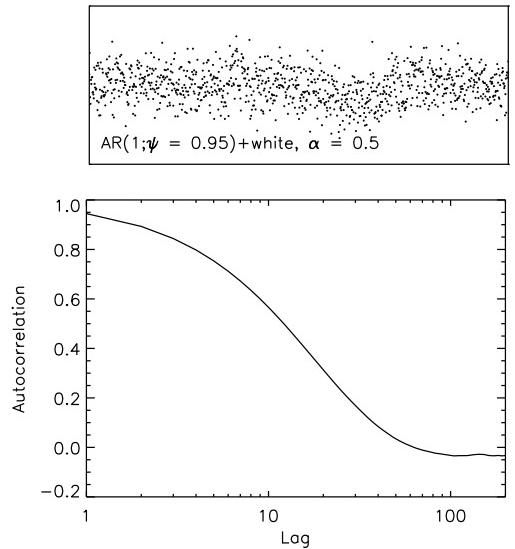


FIG. 9.— An example of an autoregressive noise process with complementary characteristics to a $1/f^\gamma$ process. The top panel shows the sum of an AR(1) process with $\psi = 0.95$ and white noise. The correlated and uncorrelated components have equal variances ($\alpha = 0.5$).

4.6. Transit timing variations estimated from a collection of light curves

We present here an illustrative calculation that is relevant to the goal of detecting planets or satellites through the perturbations they produce on the sequence of midtransit times of a known transiting planet. Typically an observer would fit the midtransit times $t_{c,i}$, to a model in which the transits are strictly periodic:

$$t_{c,i} = t_{c,0} + E_i P \quad (39)$$

for some integers E_i and constants $t_{c,0}$ and P . Then, the residuals would be computed by subtracting the best-fit model from the data, and a test for anomalies would be performed by assessing the likelihood of obtaining those residuals if the linear model were correct. Assuming there are N data points with normally-distributed, independent errors, the likelihood is given by a χ^2 -distribution, $prob(\chi^2, N_{dof})$, where

$$\chi^2 = \sum_i \left[\frac{t_{c,i} - (t_{c,0} + E_i P)}{\sigma_{t_{c,i}}} \right]^2 \quad (40)$$

and $N_{dof} = N - 2$ is the number of degrees of freedom. Values of χ^2 with a low probability of occurrence indicate the linear

TABLE 5
ESTIMATES OF t_c FROM DATA WITH AUTOREGRESSIVE CORRELATED NOISE

Method	$\langle \hat{\sigma}_{t_c} \rangle$ [sec]	$\langle \mathcal{N} \rangle$	$\sigma_{\mathcal{N}}$	$prob(\mathcal{N} > 1)$	$prob(\text{better})^a$
White	4.5	-0.010	2.50	70%	—
Wavelet	8.7	-0.016	1.33	44%	51%
Time-averaging	9.9	-0.010	1.25	40%	49%
Residual-permutation	10.2	-0.010	1.23	38%	51%

^a The probability that the analysis method returns an estimate of t_c that is closer to the true value than the white analysis.

TABLE 6
LINEAR FITS TO ESTIMATED MIDTRANSIT TIMES

Method	Fitted Period / True Period	χ^2/N_{dof}	$prob(\chi^2 < \hat{\chi}^2)$
White	$1.00000071 \pm 0.00000043$	2.25	98%
Wavelet	$1.00000048 \pm 0.00000077$	0.93	51%

model is deficient, that there are significant anomalies in the timing data, and that further observations are warranted.

We produced 10 simulated light curves of transits of the particular planet GJ 436b, a Neptune-sized planet transiting an M dwarf (Butler et al. 2004, Gillon et al. 2007) which has been the subject of several transit-timing studies (see, e.g., Ribas et al. 2008, Alonso et al. 2008, Coughlin et al. 2008). Our chosen parameters were $R_p/R_\star = 0.084$, $a/R_\star = 12.25$, $i = 85.94$ deg, and $P = 2.644$ d. This gives $\delta = 0.007$, $T = 1$ hr, and $\tau = 0.24$ hr. We chose limb-darkening parameters as appropriate for the SDSS r band (Claret 2004). We assumed that 10 consecutive transits were observed, in each case giving 512 uniformly-sampled flux measurements over 2.5 hours centered on the transit time. Noise was synthesized in the Fourier domain (as in § 4.2), with a white component $\sigma_w = 0.001$ and a $1/f$ component with rms 0.0005 ($\alpha = 1/2$). The 10 simulated light curves are plotted in Fig. (10). Visually, they resemble the best light curves that have been obtained for this system.

To estimate the midtransit time of each simulated light curve, we performed a wavelet analysis and a white analysis, allowing only the midtransit time and the noise parameters to vary while fixing the other parameter values at their true values. We used the same MCMC technique that was described in § 4.2. Each analysis method produced a collection of 10 midtransit times and error bars. These 10 data points were then fitted to the linear model of Eqn. (39). Fig. (11) shows the residuals of the linear fit (observed – calculated). Table 6 gives the best-fit period for each analysis (wavelet or white), along with the associated values of χ^2 .

As was expected from the results of § 4.2, the white analysis gave error bars that are too small, particularly for epochs 4 and 7. As a result, the practitioner of the white analysis would have rejected the hypothesis of a constant orbital period with 98% confidence. In addition, the white analysis gave an estimate for the orbital period that is more than 1σ away from the true value, which might have complicated the planning and execution of future observations. The wavelet method, in contrast, neither underestimated nor overestimated the errors, giving $\chi^2 \approx N_{\text{dof}}$ in excellent agreement with the hypothesis of a constant orbital period. The wavelet method also gave an estimate for the orbital period within 1σ of the true value.

4.7. Estimation of multiple parameters

Thus far we have focused exclusively on the determination of the midtransit time, in the interest of simplicity. However, there is no obstacle to using the wavelet method to estimate multiple parameters, even when there are strong degeneracies among them. In this section we test and illustrate the ability of the wavelet method to solve for all the parameters of a transit light curve, along with the noise parameters.

We modeled the transit as in §§ 4.1 and 4.2. The noise was synthesized in the frequency domain (as in § 4.2), using $\sigma_w = 0.0045$, $\gamma = 1$, and $\alpha = 1/2$. The resulting simulated light curve is the upper time series in Fig. 12. We used the MCMC method to estimate the transit parameters $\{R_p/R_\star, a/R_\star, i, t_c\}$ and the noise parameters $\{\sigma_r, \sigma_w\}$ (again fixing $\gamma = 1$ for simplicity). The likelihood was evaluated with either the wavelet method [Eqn. (32)] or the white method [Eqn. (6)].

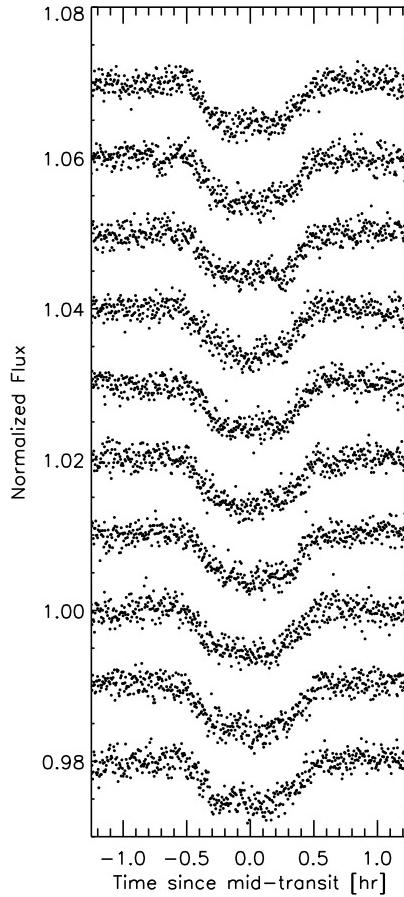


FIG. 10.—Simulated transit observations of the “Hot Neptune” GJ 436. Arbitrary vertical offsets have been applied to the light curves, to separate them on the page.

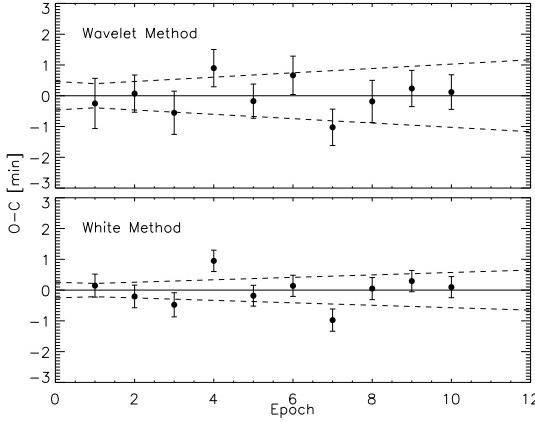


FIG. 11.—Transit timing variations estimated from simulated transit observations of GJ 436b. Each panel shows the residuals (observed – calculated) of a linear fit to the estimated midtransit times. The midtransit times were estimated with a wavelet analysis and also with a white analysis, as described in the text. The dashed lines indicate the 1σ errors in the linear model.

Fig. 13 displays the results of this analysis in the form of the posterior distribution for the case of t_c , and the joint posterior confidence regions for the other cases. The wavelet method gives larger (and more appropriate) confidence regions than the white analysis. In accordance with our previous findings, the white analysis underestimates the error in t_c and gives an estimate of t_c that differs from the true value by more than 1σ . The wavelet method gives better agreement. Both analyses give an estimate for R_p/R_\star that is smaller than the true value of 0.15, but in the case of the white analysis, this shift is deemed significant, thereby ruling out the correct answer with more than 95% confidence. In the wavelet analysis, the true value of R_p/R_\star is well within the 68% confidence region. Both the wavelet and white analyses give accurate values of a/R_\star and the inclination, and the wavelet method reports larger errors. As shown in Fig. (13), the wavelet method was successful at identifying the parameters (α and σ_w) of the underlying $1/f$ noise process.

Fig. 12 shows the best-fitting transit model, and also illustrates the action of the “whitening” filter that was described in § 3.4. The jagged line plotted over the upper time series is the best estimate of the $1/f$ contribution to the noise, found by applying the whitening filter [Eqn. (29)] to the data using the estimated noise parameters. The lower time series is the whitened data, in which the $1/f$ component has been subtracted. Finally, in Fig. 14 we compare the estimated $1/f$ noise component with the actual $1/f$ component used to generate the data. Possibly, by isolating the correlated component in this way, and investigating its relation to other observable parameters, the physical origin of the noise could be identified and understood.

5. SUMMARY AND DISCUSSION

In this paper we have introduced a technique for parameter estimation based on fitting a parametric model to a time series that may be contaminated by temporally correlated noise with a $1/f^\gamma$ power spectral density. The essence of the technique is to calculate the likelihood function in a wavelet basis. This is advantageous because a broad class of realistic noise processes produce a nearly diagonal covariance matrix in the wavelet basis, and because fast methods for computing wavelet transforms are available. We have tested and illus-

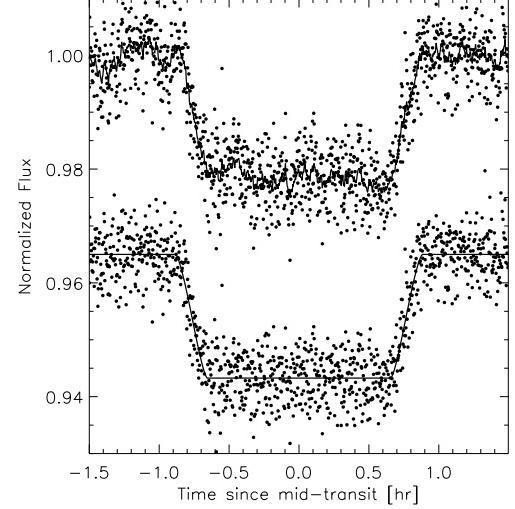


FIG. 12.—Wavelet analysis of a single simulated transit light curve. *Top.*—Simulated light curve with correlated noise. The jagged line is the best-fitting transit model plus the best-fitting model of the $1/f$ component of the noise. *Bottom.*—Simulated light curve after applying the whitening filter of Eqn. (29), using the noise parameters estimated from the wavelet analysis. The solid line is the best-fitting transit model.

trated this technique, and compared it to other techniques, using numerical experiments involving simulated photometric observations of exoplanetary transits.

For convenience we summarize the likelihood calculation here:

- Given the N data points $y(t_i)$ obtained at evenly-spaced times t_i , subtract the model $f_i(t_i; \bar{p})$ with model parameters \bar{p} to form the N residuals $r(t_i) \equiv y(t_i) - f(t_i; \bar{p})$.
- If N is not a multiple of a power of two, either truncate the time series or enlarge it by padding it with zeros, until $N = n_0 2^M$ for some $n_0 > 0, M > 0$.
- Apply the Fast Wavelet Transform (FWT) to the residuals to obtain $n_0(2^M - 1)$ wavelet coefficients r_n^m and n_0 scaling coefficients \bar{r}_n^l .
- For stationary, Gaussian noise built from an additive combination of uncorrelated and correlated noise (with Power Spectral Density $S(f) \propto 1/f^\gamma$), the likelihood for the residuals $r(t_i)$ is given by

$$\mathcal{L} = \left\{ \prod_{m=2}^M \prod_{n=1}^{n_0 2^{m-1}} \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp \left[-\frac{(r_n^m)^2}{2\sigma_W^2} \right] \right\} \times \left\{ \prod_{n=1}^{n_0} \frac{1}{\sqrt{2\pi\sigma_S^2}} \exp \left[-\frac{(\bar{r}_n^l)^2}{2\sigma_S^2} \right] \right\} \quad (41)$$

where

$$\sigma_W^2 = \sigma_r^2 2^{-\gamma m} + \sigma_w^2 \quad (42)$$

$$\sigma_S^2 = \sigma_r^2 2^{-\gamma} g(\gamma) + \sigma_w^2 \quad (43)$$

for some noise parameters $\sigma_w > 0, \sigma_r > 0$ and $g(\gamma) = O(1)$ [e.g., $g(1) \approx 0.72$].

The calculation entails the multiplication of N terms and has an overall time-complexity of $O(N)$. With this prescription

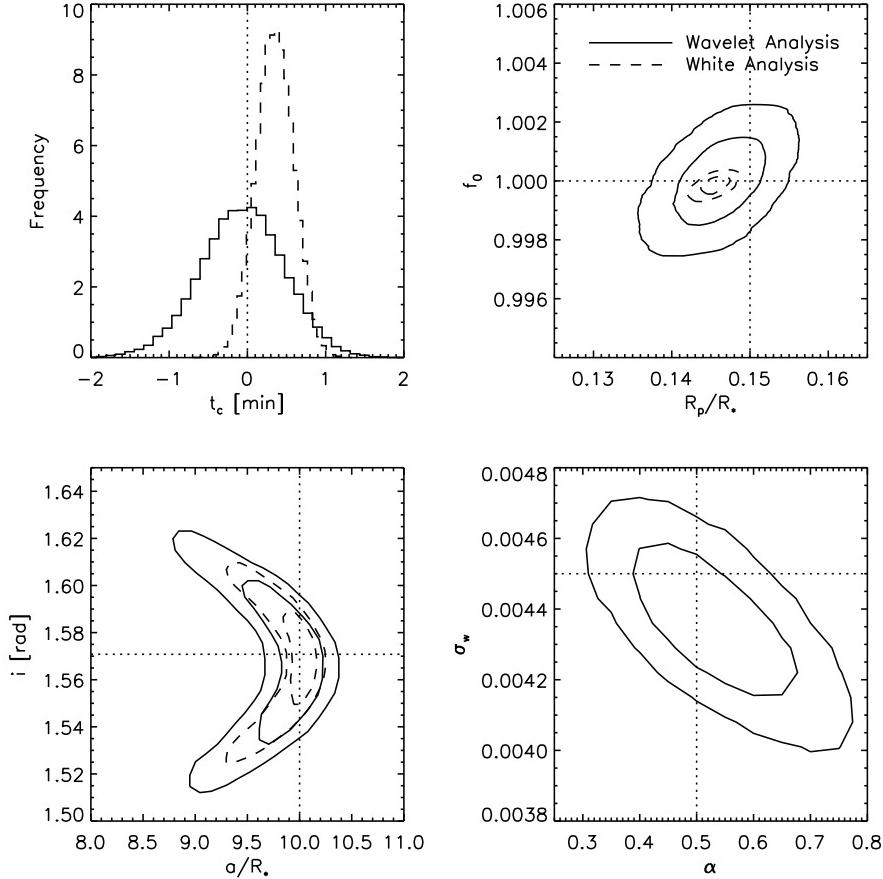


FIG. 13.— Results of parameter estimation for the simulated light curve of Fig. 12. Results for both the wavelet method (solid lines) and the white method (dashed lines) are compared. The upper left panel shows the posterior distribution for the midtransit time. The other panels show confidence contours (68.3% and 95.4%) of the joint posterior distribution of two parameters. The true parameter values are indicated by dotted lines.

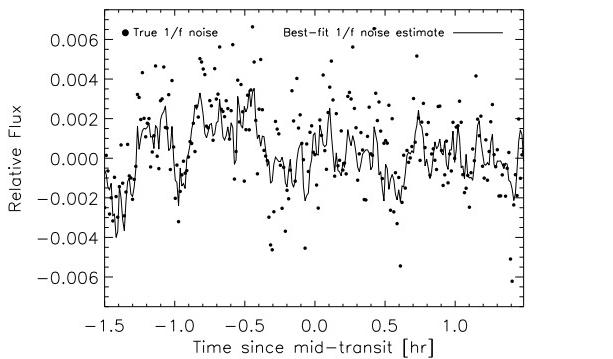


FIG. 14.— Isolating the correlated component. Plotted are the actual and estimated $1/f$ components of the noise in the simulated light curve plotted in Fig. (12). The estimated $1/f$ signal was found by applying the wavelet filter, Eqn. (29), to the residuals.

for the likelihood function, the parameters may be optimized using any number of traditional algorithms. For example, the likelihood may be used in the jump-transition probability in a Monte Carlo Markov Chain analysis, as we have done in this work.

Among the premises of this technique are that the correlations among the wavelet and scaling coefficients are small enough to be negligible. In fact, the magnitude of the correlations at different scales and times are dependent on the choice of wavelet basis and the spectral index γ describing

the power spectral density of the correlated component of the noise. We have chosen for our experiments the Daubechies 4th-order wavelet basis which seems well-suited to the cases we considered. A perhaps more serious limitation is that the noise should be stationary. Real noise is often nonstationary. For example, photometric observations are noisier during periods of poor weather, and even in good conditions there may be more noise at the beginning or end of the night when the target is observed through the largest airmass. It is possible that this limitation could be overcome with more elaborate noise models, or by analyzing the time series in separate segments; future work on these topics may be warranted.

Apart from the utility of the wavelet method, we draw the following conclusions based on the numerical experiments of § 3. First, any analysis that ignores possible correlated errors (a “white” analysis in our terminology) is suspect, and any $2-3\sigma$ results from such an analysis should be regarded as provisional at best. As shown in §§ 4.1, 4.2, and 4.6, even data that appear “good” on visual inspection and that are dominated by uncorrelated noise may give parameter errors that are underestimated by a factor of 2–3 in a white analysis. Second, using any of the methods described in 4.4 (the wavelet method, the time-averaging method, or the residual-permutation method) is preferable to ignoring correlated noise altogether.

Throughout this work our main application has been estimation of the parameters of a single time series or a few such time series, especially determining the midtransit times of transit light curves. One potentially important application

that we have not discussed is the *detection* of transits in a database of time-series photometry of many stars. Photometric surveys such as the ground-based HAT (Bakos et al. 2007) and SuperWASP (Pollacco et al. 2006), and space-based missions such as Corot (Baglin et al. 2003) and Kepler (Borucki et al. 2003) produce tens to hundreds of thousands of time series, spanning much longer intervals than the transit durations. It seems likely that the parameters of a noise model could be very well constrained using these vast databases, and that the application of a wavelet-based whitening filter could facilitate the detection of transits and the elimination of statistical false positives. Popular techniques for dealing with correlated

noise in large photometric databases are those of Tamuz et al. (2005), Kovács et al. (2005), and Pont et al. (2006). A priority for future work is to compare these methods with a wavelet-based method, by experimenting with realistic survey data.

We are grateful to Frederic Pont and Eric Feigelson for very detailed and constructive critiques of an early version of this manuscript. We also thank Scott Gaudi and Jason Eastman for helpful comments. This work was partly supported by the NASA Origins program (grant no. NNX09AB33G).

REFERENCES

- Agol, E., Steffen, J., Sari, R., & Clarkson, W. 2005, MNRAS, 359, 567
 Alonso, R., Barbieri, M., Rabus, M., Deeg, H. J., Belmonte, J. A., & Almenara, J. M. 2008, A&A, 487, L5
 Baglin, A. and the COROT team 2003, Advances in Space Research, 31, 345
 Bakos, G. Á., et al. 2006, ApJ, 650, 1160
 Bakos, G. Á., et al. 2007, ApJ, 656, 552
 Bean, J. L., et al. 2008, A&A, 486, 1039
 Beran, J. 1994, Statistics for Long-Memory Processes (London: Chapman & Hall)
 Bevington, P. R., & Robinson, D. K. 2003, Data reduction and error analysis for the physical sciences, 3rd ed. (Boston, MA: McGraw-Hill)
 Borucki, W. J., et al. 2003, Proc. SPIE, 4854, 129
 Bouchy, F., Pont, F., Melo, C., Santos, N. C., Mayor, M., Queloz, D., & Udry, S. 2005, A&A, 431, 1105
 Box, G. E. P., & Jenkins, G. M. 1976, Holden-Day Series in Time Series Analysis, Revised ed., San Francisco: Holden-Day, 1976,
 Bracewell, R. The Fourier Transform and Its Applications. McGraw-Hill Electrical and Electronic Engineering Series, Terman, F., Ed. New York: McGraw-Hill, 1965, pg. 115
 Burke, C. J., et al. 2007, ApJ, 671, 2115
 Butler, R. P., Vogt, S. S., Marcy, G. W., Fischer, D. A., Wright, J. T., Henry, G. W., Laughlin, G., & Lissauer, J. J. 2004, ApJ, 617, 580
 Carter, J. A., Yee, J. C., Eastman, J., Gaudi, B. S., & Winn, J. N. 2008, ApJ, 689, 499
 Claret, A. 2004, A&A, 428, 1001
 Collier Cameron, A., et al. 2007, MNRAS, 380, 1230
 Coughlin, J. L., Stringfellow, G. S., Becker, A. C., López-Morales, M., Mezzalira, F., & Krajci, T. 2008, ApJ, 689, L149
 Daubechies, I. 1988, Communications on Pure and Applied Mathematics, 41, 7, 909
 Fadili, M.J. & Bullmore, E.T. 2002, *Neuroimage*, 15, 217
 Gibson, N. P., et al. 2008, A&A, 492, 603
 Gillon, M., Pont, F., Moutou, C., Bouchy, F., Courbin, F., Sohy, S., & Magain, P. 2006, A&A, 459, 249
 Gillon, M., et al. 2007, A&A, 472, L13
 Gillon, M., et al. 2009, A&A, 496, 259
 Giménez, A. 2007, A&A, 474, 1049
 Gould, A. 2003, ArXiv Astrophysics e-prints, arXiv:astro-ph/0310577
 Gregory, P. C. 2005, Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with *Mathematica* Support (Cambridge Univ. Press, UK).
 Haar, A., Zur Theorie der Orthogonalen Funktionensysteme., Math. Annal., Vol. 69, pp. 331-371, 1910
 Holman, M. J., & Murray, N. W. 2005, Science, 307, 1288
 Holman, M. J., et al. 2006, ApJ, 652, 1715
 Jenkins, J. M., Caldwell, D. A., & Borucki, W. J. 2002, ApJ, 564, 495
 Johnson, J. A., Winn, J. N., Cabrera, N. E., & Carter, J. A. 2009, ApJ, 692, L100
 Kipping, D. M. 2009, MNRAS, 392, 181
 Koen, C., & Lombard, F. 1993, MNRAS, 263, 287
 Konig, M., & Timmer, J. 1997, A&AS, 124, 589
 Kovács, G., Bakos, G., & Noyes, R. W. 2005, MNRAS, 356, 557
 Mallat, S. G. 1989, IEEE Transactions on Pattern Analysis and Machine Intelligence, 11, 674
 Mallat, S. A Wavelet Tour of Signal Processing, London:AP Professional, 1997
 Mandel, K., & Agol, E. 2002, ApJ, 580, L171
 Percival D.B., & A.T. Walden. Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques. Cambridge: Cambridge University Press, 1993
 Pollacco, D. L., et al. 2006, PASP, 118, 1407
 Pont, F., Zucker, S., & Queloz, D. 2006, MNRAS, 373, 231
 Press, W. H. 1978, Comments on Astrophysics, 7, 103
 Press, William H., Teukolsky, Saul A., Vetterling, William T., & Flannery, Brian P. 2007, Numerical Recipes: The Art of Scientific Computing, 3rd ed. (New York: Cambridge University Press)
 Ribas, I., Font-Ribera, A., & Beaulieu, J.-P. 2008, ApJ, 677, L59
 Southworth, J. 2008, MNRAS, 386, 1644
 Tamuz, O., Mazeh, T., & Zucker, S. 2005, MNRAS, 356, 1466
 Teolis, A. Computational Signal Processing with Wavelets. Boston:Birkhäuser, 1998
 Timmer, J., Schwarz, U., Voss, H. U., Wardinski, I., Belloni, T., Hasinger, G., van der Klis, M., & Kurths, J. 2000, Phys. Rev. E, 61, 1342
 Udalski, A., Szymanski, M. K., Kubiak, M., Pietrzynski, G., Soszynski, I., Zebrun, K., Szewczyk, O., & Wyrzykowski, L. 2004, Acta Astronomica, 54, 313
 Winn, J. N., et al. 2007, AJ, 133, 1828
 Winn, J. N., et al. 2008, ApJ, 683, 1076
 Winn, J. N., Holman, M. J., Carter, J. A., Torres, G., Osip, D. J., & Beatty, T. 2009, AJ, 137, 3826
 Wornell, G. Signal Processing with Fractals: A Wavelet-Based Approach. Prentice Hall Signal Processing Series, Oppenheim, A., Ed. Upper Saddle River, NJ:Prentice Hall Press, 1996